

TIME VARYING FILTER WITH ZERO AND/OR POLE MIGRATION

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of United States Provisional Patent Application No. 60/317,802, filed on September 7, 2001.

TECHNICAL FIELD

This invention relates to signal processing, and more particularly, to an improved technique for varying a filter used to process a given audio signal, so as to remove audible artifacts due to such filter variance and thus improve audio quality.

BACKGROUND OF THE INVENTION

In digital signal processing applications, such as filtering of audio signals to improve audio quality, the particular filter needed to optimally filter a given input signal often varies over the course of the input signal. Thus, pole and zero locations of the various filters used with the various segments of the input signal will in general be different. Simply switching from one implemented filter to another during processing presents a problem. The filter switch results in an audio artifact, sometimes described as a "click". The click is both audible and annoying to the user.

For example, suppose in a dictation application the original dictator, dictates a long speech. His environment is such that there is background noise which changes over the course of the dictation. Suppose an airplane flies overhead at one moment, a siren blares in the background at another, a buzzer or bell due to some maintenance function at the

dictator's location goes off in a third. Such background sounds are undesirable noise relative to the information content of the dictated speech. Optimally this background noise should be deleted. In the simple case where the noise is an unwanted spectral line in the frequency characterization of the input signal, for example, a notch filter can be used to suppress such an unwanted tone. However, as stated above, if the frequency of the unwanted spectral tone, or some combination of unwanted tones, varies over different portions of the input audio signal, different filters must be used. Hence the idea of time varying filtering.

A filter in general will be desired for one segment of an input signal, but for the other segments will not be desired. Thus, filters must be switched in and out over the duration of the input signal. Alternatively, the various filters could be suppressed by switching in a bypass filter (which would waste computing resources as the "old" filters and their bypass filters accumulate, and each one continues to have its filtering implemented by the system), or the once useful filter can simply be left in place, and continue to filter segments it is useless on. In conventional techniques for the filtering of audio signals, there is really no other alternative. Either (a) a filter is switched in and out, or suppressed by a bypass filter, or (b) a filter is left in place. If the filter is switched abruptly in and out, or a bypass filter is switched in, audible artifacts are generated. If a filter is kept continually in place the recording quality is reduced during the portions of the recording session when the unwanted tone or tones, which that particular filter was designed to remove, are not

present. What is desired, therefore, is a means of implementing time varying filtering seamlessly and in a manner inaudible to the user.

It is an object of the present invention to solve the above-described problems inherent in the art, and implement a clickless, or inaudible, method of inserting and
5 removing a filter, or other digital signal processing operator.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a z-plane diagram depicting an engaged filter according to the present invention;

10 Figure 2 is a z-plane diagram depicting the filter of Figure 1 after zero migration according to the present invention;

Figure 3 is a z-plane diagram depicting migration of poles and zeros to the origin according to one embodiment of the present invention;

15 Figure 4 is a z-plane diagram depicting a second engaged filter according to the present invention;

Figure 5 is a z-plane diagram depicting the filter of Figure 4 after zero migration according to the present invention; and

Figure 6 is a state machine diagram indicating operational flow according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The idea of the present invention is a simple one. It is a clickless, i.e. without audible artifact, method of inserting and removing a filter by gradually migrating the coefficients of the filter from their original values to values wherein the filter has no effect.

Conversely, the method in audibly switching on the filter is accomplished by gradually changing the coefficients of such filter from values wherein the filter has no effect to their final values where the filter has the desired effect.

For purposes of illustration, the invention will be illustrated by the insertion and removal of a notch filter, or the switching on and off of such notch filter. Frequency domain analysis will be utilized to illustrate the filters and their engagement and disengagement. Thus, a filter will be illustrated by the locus of its poles and zeros in the z-plane. The changing of the coefficients of the filter from an active to an inactive configuration will be illustrated by migrating the poles and zeros of the notch filter from one locus of positions in the z-plane to another locus of positions. It is understood that the invention can be implemented using any alternative method of visualizing filters, both digital and continuous, and thus the example of a notch filter and the herein described embodiment utilizing the z-plane are intended solely for illustrative purposes and are in no way intended to limit the invention or its various embodiments.

Figure 1 depicts a fourth order notch filter in the z-plane. As can be seen, there are two complex conjugate pole pairs comprising P1 and P3 in the first quadrant, and their

corresponding poles P2 and P4 in the fourth quadrant. As well, there are two zeros collocated at points Z1, Z3 in the first quadrant and Z2, Z4 in the fourth quadrant. This notch filter removes the frequency or unwanted spectral line at $\pi/4$ radians (i.e., $F_s/8$ Hz, where F_s is the sampling frequency). This configuration of the poles and zeros will filter out the frequency of $\pi/4$ radians (i.e., $F_s/8$ Hz). In order to disable this filter, i.e., change the pole and zero locations so that it has no effect whatsoever, is a simple matter. All that needs to be done is to migrate the zeros from their location on the unit circle at $\pm \pi/4$ (i.e., $\pm F_s/8$ Hz) to the locations of the poles, such that there remain four points in the z-plane, each with a pole and a canceling zero. For this to be done abruptly, it would create an audible artifact in an audio signal, and some type of anomaly in a different signal type. The method of invention is to do it gradually, in a series of intermediate steps.

The more gradually the filter is disengaged the larger the segment of the input signal that is involved. Conversely the less gradually, or the more quickly, that the filter is disengaged, fewer samples will be involved. The number of samples per step in the migration of the zeros to the pole locations in general is found empirically for each filter and each signal in each general type of input signal. For the depicted embodiment shown in Figure 1, the zeros are migrated to the pole locations in 300 intermediate steps, with one input sample per step. Thus, for each subsequent sample, each zero is moved to a different intermediate location until it reaches its corresponding pole. Should a more gradual (or abrupt) disengagement be desired the number of steps could be increased (or

decreased), and thus the incremental change of coefficient values from step to step decreased (or increased). As well, numerous samples could be processed at each intermediate filter step as opposed to the one sample per step implementation. Figure 2 depicts the completely disengaged notch filter showing the four zeros each collocational with one of the poles. Figure 2 therefore depicts points P1, P2, P3, and P4 where at each such point there is a pole and a zero. These four locations correlate exactly to the original pole locations P1, P2, P3 and P4, respectively, depicted in Figure 1.

The locus of the poles and zeros depicted in Figure 2 results in a filter which has no effect, and is thus wholly disengaged.

Once the poles and zeros are collocational, although the effect of the filter is neutralized, since the filter is still operating – although with no net effect -- extra mathematical operations are still being performed. These operations cause reduced numerical accuracy and increased processor workload, expressed in Millions of Instructions per Second, or “MIPS.” Thus, in one embodiment of the invention, after the state depicted in Figure 2 has been reached, the pole/zero pairs comprising the now neutralized filter are removed. This may be done in two ways. The filter can simply be removed, or bypassed, after the collocational points have been reached. Alternatively, the collocational points can be migrated to the origin. Once at the origin, the filter can be abruptly removed without any artifacts. In either case, removal of the filter reduces the MIPS usage, which may be of great utility in dynamic systems, which allocate MIPS

adaptively, such as multichannel systems.

Figure 3 depicts the migration to the origin method, showing the four pole/zero colocal points, identical to those depicted in Figure 2, now being migrated to the origin for removal of the entire filter upon reaching the origin. During the entire migration of the pole/zero pairs from their final locations, as depicted in Figure 2, to the origin, the filter remains neutralized and has no effect on the signal.

In general, the number of steps used in the migration of the zeros to the poles, or vice versa, as the case may be, in neutralizing a filter, will be determined empirically by the user. As well, the same determination as to the number of steps will be made by the user if after neutralization the colocal points are migrated to the origin, and from there the filter bypassed. The operational tradeoffs in determining the number of such steps are speed of removal of an unwanted filter vs. tolerable audible artifacts.

As well, whether to bypass the filter, as in the above described embodiment, or simply to leave it operating with no net effect will depend upon the application. In some applications, it may not be necessary to turn off the filter after neutralization. Although this can cause reduced accuracy and consumes more MIPS, in some applications MIPS availability is static, designed for the worst case MIPS requirements. Thus, there is no benefit realized by turning off the filter. Accuracy is also not always an issue, as real-world noise and distortion in a given environment can greatly outweigh the precision of the digital signal processor ("DSP"). In such applications, simplicity and less

control code are more important than wasting already statically allocated MIPS.

On the other hand, in an application where the DSP been designed to operate on a large number of channels, and if it had software to take advantage of the statistical nature of the application (a resource manager and a way to degrade slowly rather than hard fail), turning off the filter can be a very good idea. Turning off the filter may be very important in other applications, too.

The theoretical and practical issues involved in turning off the filter after neutralization are as follows. For simplicity, consider a first order system with a single zero at $z=b$ and a single pole at $z=a$: $H(z) = (z-b)/(z-a)$. It is assumed, for illustration purposes, that the zero is migrated and the pole left fixed (as would be the normal case with a notch filter, as described above). Assume further, that in this illustration the notch occurs at a frequency of 0. The time domain equation is: $y[k] = a*y[k-1] + u[k] - b*u[k-1]$. Again, for illustrative purposes only, consider that the zero moved instantly at some point in time from $b=1$ to $b=a$. Now, say that at some subsequent arbitrary time, the filter is to be turned off, i.e., $y[k] = u[k]$. In theory, if the filter has been excited (as would be the normal case for a non-zero input signal), $a*y[k-1]$ will **not** equal $a*u[k-1]$. The two are only equal **after** the filter transient has settled. If the coincident zero/pole pair is migrated toward zero, this is effectively forcing the terms $a*y[k-1]$ and $a*u[k-1]$ to zero. Once they are zero, the filter can be bypassed. This describes the theoretical model for an abrupt change

in the zero from $b=1$ to $b=a$.

Practically, however, the case would generally be different.

The zero would generally be slowly migrated to the pole position. If this occurs slowly

enough, then the transient would most likely (except for a nearly unstable

filter) be decayed to a small value. Thus, turning off the filter at the collocational pole zero

points would result in little or no audible glitch, and would be the preferred solution. If there

was an audible glitch, what could be done is to simply wait

a bit (this will allow the transient to decay) and then bypass the filter.

It is noted that migrating the pole and zeros toward the origin makes the filter

more stable, and thus speeds the decay of the transients.

Thus far, the gradual, and thus inaudible, disengaging of a filter has been described.

The inverse of neutralizing a filter is engaging, or enabling, a filter. For purposes of

illustration, it is assumed that the unwanted spectral line in the signal frequency spectrum,

or unwanted tone, at $\pi/4$ (i.e., $F_s/8$ Hz) which was bothersome in the first part of the signal

(as described above) is no longer present in a second part of the signal. However, the

second part of the signal has an annoying unwanted spectral line at $3\pi/8$ radians (i.e.,

$F_s \cdot 3/16$ Hz). Thus, to suppress the unwanted spectral line at $3\pi/8$ radians (i.e., $F_s \cdot 3/16$

Hz), the filter depicted in Figure 4 is used. Once again, this is a fourth order notch filter

designed to remove the unwanted tone in the second segment of the input signal, now at

$3\pi/8$ (i.e., $F_s \cdot 3/16$ Hz). As can be seen, two complex conjugate pairs of poles and two

colocational pairs of zeros at each of $\pm 3\pi/8$ (i.e., $\pm F_s \cdot 3/16$ Hz) comprise the filter. For the reasons discussed above, this filter can simply not just be turned on but it needs to be gradually engaged so as to avoid audible artifacts in the filtered, or output signal. Utilizing exactly the inverse of the neutralizing method of the filter depicted in Figures 1 and 2, to engage the filter of Figure 4, there is first implemented the filter in Figure 5 which is nothing more than the notch filter depicted in Figure 4 with the zeros now colocational with the four poles, thus being a neutralized filter. To engage this filter, the zeros are migrated over a series of intermediate locations, along the dotted lines 410, to their final positions on the unit circle. In all the figures, the zeros are migrated, whether from the unit circle, or to the unit circle, along the shortest line in the z-plane between their original locations and their final locations.

In general, a filter will have various poles and zeros, and either can be migrated to the other to reach a destination point of each pole with its corresponding zero. The migration should be such as to most efficiently disengage the filter. Thus, the shortest path between the poles and zeros should be utilized for the migration. The particular manner in which the poles and zeros are migrated to each other can be optimized for the filter being used, the characteristics of the noise being filtered, as well as the characteristics of the input signal being operated upon. Thus, in the filter depicted in Figure 1, since any migration of the poles toward the zeros would tend to increase the effect of the poles, and leave the zeros at their full influence, the zeros are migrated towards the poles, along the

shortest path between them, as indicated by the dotted lines in Fig. 1.

In this embodiment each pole zero pair, or other final configuration of the filter, will then be removed, to reduce wasted computational resources. As described above, generally the filter will be bypassed after neutralization, but may, if necessary, have the colocal pole zero pairs first migrated to the origin.

Figure 6 depicts a state machine diagram for the example filter depicted in Fig. 1, where the filter is bypassed, but never removed (thus neither removal of the filter from the final configuration shown in Fig. 3, nor the removal of the filter from the configuration of Figure 2, occurs in this example system). The binary variable NotchRequest is the input to the state machine. As described above, each particular application, as well as the subjective perception of the users interacting with the output signal, will determine the minimum steps needed to fall below the threshold of perceptible artifacts (auditory or otherwise, as determined by the signal being processed). Such step numbers can be empirically found with relative simplicity. The number of steps used is the variable "STEPS" to a coefficient calculation function, labeled Calc Coeff in Fig. 6. In the system of Fig. 6, the Calc Coeff function is determined by parameterizing the migration paths, and STEPS=300. Consider the migration of the zeros of the depicted fourth order notch filter of Figure 1. In this case, for a notch filter, the poles and zeros align on straight lines in the z-plane, which are the dotted lines in Figure 1. The co-efficients can be calculated in real time, or if the filter used in the system is known *a priori*, as well as the desired step

number, the various coefficients for each step can be stored in a lookup table, moving thereby the complex high MIPS calculations to non-real-time at the expense of a small to moderate table.

In Figure 6 the state 600 "Notch Bypass", is that where no filter is operating, and the poles and zeros are colocal, as in Fig. 2. In state 620 "Notch ON" the filter is fully engaged. States 610 and 630 represent the migration stages "Migrate Notch IN" and "Migrate Notch OUT." In state 610 the filter is migrated from a neutralized state to the fully engaged state, via the intermediate steps. In state 630 filter is migrated from a fully engaged state, to the neutralized state, via the intermediate steps. The number of intermediate steps is controlled by the variable "STEPS."

The system moves from state 600 to state 610, along path 601, when Notch Request is set to ON. In state 610 the filter is migrated from being bypassed (i.e., the poles and zeros are colocal), as depicted in Figure 2, through the 300 intermediate steps, to being fully implemented, as depicted in Figure 1. Once the final step has been reached, i.e., STEPS=300 in the illustrative system of Fig. 6, the state moves, along path 611, to state 620, "Notch ON." From there a Notch Request = OFF will move the system, along path 621, to state 630, where the reverse happens, and the poles and zeros in the configuration of Fig. 1 now migrate to the configuration of Figure 2, via the 300 intermediate steps. When, in state 630, STEPS =300, the filter is totally migrated OUT, and the system moves, along path 631, to state 600.

As well, if in the middle of a filter migration one way or the other the Notch Request value changes, the system will move from state 610 to state 630 or vice versa. With reference to Figure 6, if in state 630, where the filter is migrating OUT due to a Notch Request = OFF signal, Notch Request switches to ON, the system moves, along path 632, to state 610, and migrates the filter back IN, beginning from whatever the value of STEP was when the Notch Request signal changed from OFF to ON. Conversely, if in the middle of a migration IN, where the system is in state 610, Notch Request goes to OFF, the system moves, along path 612 to state 630, and the migration OUT proceeds from whatever the value of STEP was in state 610.

Although the invention has been illustrated via the z-plane depiction of filters, the actual coefficients used to operate upon the input signal are calculated from these pole and zero locations, and could be calculated from any equivalent representation of a filter, whether time domain or frequency domain. As described above, the entire invention may be implemented using a computer or other data processing devices, with the appropriately written software. The various functionalities of the invention may be implemented using hardware, software, or combinations of both, resulting in various permutations of the presently depicted embodiment.

While the foregoing describes one embodiment of the invention, it will be appreciated by those of skill in the art that various modifications and additions may be made. Such additions and modifications would include, without limitation,

5

[illegible]